

YFitter: Maximum likelihood assignment of Y chromosome haplogroups from low-coverage sequence data

Luke Jostins, Yali Xu, Shane McCarthy, Qasim Ayub,
Richard Durbin, Jeff Barrett, Chris Tyler-Smith

July 31, 2014

Abstract

Low-coverage short-read resequencing experiments have the potential to expand our understanding of Y chromosome haplogroups. However, the uncertainty associated with these experiments mean that haplogroups must be assigned probabilistically to avoid false inferences. We propose an efficient dynamic programming algorithm that can assign haplogroups by maximum likelihood, and represent the uncertainty in assignment. We apply this to both genotype and low-coverage sequencing data, and show that it can assign haplogroups accurately and with high resolution. The method is implemented as the program YFitter, which can be downloaded from <http://sourceforge.net/projects/yfitter/>

1 Introduction

Low-coverage, short read resequencing is a cost effective means of carrying out variant discovery, disease association and population genetics experiments^{5,9}. One potential value of low-coverage sequencing experiments is that, as the experiments are whole-genome, previously less well-studied regions of the genome such as the Y chromosome are sequenced “for free”. Many Y chromosome mutations have been discovered⁷, and the haplogroups they define have been found to be associated with various population genetic and disease associations^{11,12}. Large, low coverage sequencing projects have the potential to greatly refine our understanding of these Y haplogroups.

However, the uncertainty associated with the lower coverage and higher error rate of these experiments has to be handled statistically to avoid biases, and this is especially true of the Y chromosome, due to the lower sequence coverage. Missing or uncertain data can result in incorrect haplogroup assignment, especially if present high up in the haplogroup tree, which can then lead to false inferences. In addition, assigning samples by hand in large sequencing experiments can be time consuming, so an automated solution is required. While probabilistic, automatic methods have been developed for Short-Tandem Repeats (Y-STRs)², no equivalent method has been developed for low coverage sequencing.

The large number of different Y haplogroups in a given tree makes separate calculation of the likelihood of sequencing reads given each Y haplogroup computationally expensive. We propose an efficient dynamic programming algorithm to calculate the likelihood, and use this to assign maximum-likelihood haplogroups robustly. This method can accurately assign individuals to haplogroups given either genotyping chip or low-coverage sequence data, and can calculate confidence haplogroups that take into account uncertainty using the Akaike information criteria.

We have implemented this method in C++ as the program YFitter, which is open source and freely available.

2 The Method

As there is no homologous recombination on the Y chromosome, Y haplogroups lie on a tree with each node i being defined by one or more mutations B_i . Our aim is to select the Y chromosome haplogroup that maximises the likelihood of the observed reads, along with a lower resolution confidence haplogroup that encompasses all plausible haplogroups. We will do this by defining the likelihood in terms of recursively calculable statistics, using the tree structure of the haplogroups.

We will write B_i^+ if a set of mutations is present in the individual under consideration, and B_i^- if it is not. The raw data for the algorithm is the per-site likelihood of observing reads at the mutation sites D_i given that the mutations B_i have or have not occurred:

$$M_i^+ = P(D_i|B_i^+) \quad (1)$$

$$M_i^- = P(D_i|B_i^-) \quad (2)$$

Many methods have been developed to calculate such likelihoods from short read data⁴⁸.

We will define all nodes downstream of node i as $i \downarrow$, and the reads and mutations at these sites $D_{i \downarrow}$ and $B_{i \downarrow}$. We can then define the downward likelihood of node i as

$$L_i^\downarrow = P(D_{i \downarrow}|B_i^-, B_{i \downarrow}^-) \quad (3)$$

$$= \begin{cases} M_i^- & \text{for leaf nodes} \\ M_i^- \prod_{j \in d(i)} L_j^\downarrow & \text{for non-leaf nodes} \end{cases} \quad (4)$$

Where d_i is the set of daughters of node i . This is the likelihood of observing reads downstream of i , given that the individual's haplogroup assignment is not descended from i . We calculate this starting at the leaf nodes, and work upwards to the root node.

For each branch, we can then define the upwards likelihood:

$$L_i^\uparrow = P(D_{i \downarrow}|B_i^+, B_{i \uparrow}^+, B_{s(i)}^-, B_{s(i) \downarrow}^-, B_{c(i)}^-) \quad (5)$$

$$= \begin{cases} M_i^+ & \text{for the root node} \\ L_{p(i)}^\uparrow M_i^+ \prod_{j \in s(i)} L_j^\downarrow & \text{for non-root nodes} \end{cases} \quad (6)$$

Where $p(i)$ is the parent of node i , $s(i)$ is the set of siblings of i , $i \uparrow$ is the set of direct ancestors of i and $c(i)$ (for “cousins”) is the set of all nodes that are not descended from or direct ancestors of $p(i)$.

The upwards likelihood is thus the likelihood of observing reads at mutation sites not descended from i , given that the individual’s haplogroup is descended from i . This value is calculated from the root node, working down to the leaf nodes.

For leaf nodes, the full likelihood is equal to the upwards likelihood, as there is no data downstream of them. For non-leaf haplogroups, we define the likelihood as the maximum of the likelihoods of its descendants.

$$L_i = P(D|B_i^+, B_{i\uparrow}^+, B_{i\downarrow}^+, B_{s(i)}^-, B_{s(i)\downarrow}^-, B_{c(i)}^-) \quad (7)$$

$$= \begin{cases} L_i^\uparrow & \text{for leaf nodes} \\ \max(L_j : j \in d(i)) & \text{for non-leaf nodes} \end{cases} \quad (8)$$

As it is likely that multiple haplogroups will all have the maximum likelihood, the maximum likelihood haplogroup is then defined as the haplogroup with the maximum likelihood that is closest to the root node. This is equivalent to the most recent common ancestor of all haplogroups with the maximum likelihood.

As well as a maximum-likelihood haplogroup, we define a haplogroup confidence set as all haplogroups with a likelihoods within 8.685 phred-scaled log units of the maximum likelihood, equivalent to a ΔAIC of 4³. The confidence haplogroup is thus the most recent common ancestor of all haplogroups in this confidence set, and all haplogroups not derived from the confidence haplogroup are judged to have “considerably less support”³.

The YFitter program reads in a haplogroup tree in phyloXML format⁶, with mutations specified as properties of branches. We constructed such a haplogroup tree using the mutations catalogued by Karafet *et al*⁷. We removed G/C and A/T SNPs to avoid stranding errors, as well as indels, repeats and non-uniquely mapped variants.

3 Applications

3.1 Assigning Haplogroups to Genotype Data

We tested our method on publicly available data from the 9 males of the Genomes Unzipped project. All participants in the project have released genotyped data generated by the personal genomics company 23andMe. This is a good test set for haplogrouping, as the custom 23andMe chip is designed to contain a large number of haplogroup-informative variants.

We assigned haplotypes to the 9 individuals using our YFitter program, and compared the results to the assignments made by 23andMe (Table 1). The set contains 3 different major haplogroups. All individuals have broadly consistent assignments, though there is some ambiguity between the haplogroup names within haplogroup R1b1b2 between 23andMe (who use the ISOGG2010 tree) and Karafet *et al* (the YCC2008 tree).

Individual	YFitter haplogroup	23andMe haplogroup
CAA001	R1a1	R1a1a*
DBV001	J2	J2
DFC001	R1b1b2g	R1b1b2a1a1d*†
DGM001	R1a1	R1a1a*
JCB001	R1b1b2	R1b1b2a1
JKP001	R1b1b2	R1b1b2a1a2f
JXA001	R1b1b2g	R1b1b2a1a1*†
LXJ001	N1c1	N1c1
VXP001	R1b1b2d	R1b1b2a1 †

Table 1: Haplogroup assignments using YFitter, compared to 23andMe’s assignments, for the Genomes Unzipped males. Variants with inconsistent nomenclature are marked with a †.

3.2 Assigning Haplogroups to Low Coverage Sequencing Data

We also tested our method using publicly available data from 286 individuals present both in the Phase I of the 1000 Genomes Project⁵ and Phase 3 of the HapMap project¹. Haplogroups were assigned manually using the HapMap genotyping data, and automatically using YFitter on the low-coverage 1000 Genomes Project sequencing data. Genotype likelihoods were generated from sequence data using the program samtools⁸. The 286 individuals contained 12 different major haplogroups.

Of the 286 maximum likelihood haplogroup assignments, 285 were fully consistent between the genotype and sequence data. Of those, 203 assignments had greater resolution in the sequence data, 71 had the same resolution, and in 11 had a lower resolution. If the confidence haplogroup was used, there were no inconsistencies, 199 had higher resolution, 75 had the same resolution, and 12 had lower resolution. Both sets sequenced-based haplogroup assignments were of higher resolution than the genotype-based set, and the confidence haplogroups were only of slightly lower resolution than the maximum likelihood haplogroups (Figure 1).

4 Discussion

We have presented an efficient statistical method for assigning haplogroups by maximum likelihood, and shown that it can accurately and automatically assign haplogroups to short-read data.

YFitter can be downloaded from <http://sourceforge.net/projects/yfitter/>. The software can either be using in conjunction with the short-read utility program samtools⁸ to assign haplogroups to sequencing data, or with the genotype utility program PLINK¹⁰ to assign haplogroups to genotype data.




Figure 1: The distribution of haplogroup depth across the manually assignment samples using HapMap data, and the confidence and maximum likelihood assignments from the low coverage data. A haplogroup depth of zero represents samples that cannot be assigned to any major haplogroup, a depth of one represents assignment to a major haplogroup but no further resolution, and each additional assignment adds one to the haplogroup depth.

References

- [1] D.M. Altshuler, R.A. Gibbs, L. Peltonen, S.F. Schaffner, F. Yu, E. Dermitzakis, P.E. Bonnen, P.I.W. de Bakker, P. Deloukas, S.B. Gabriel, et al. Integrating common and rare genetic variation in diverse human populations. *Nature*, 467:52–58, 2010.
- [2] T.W. Athey. Haplogroup prediction from y-str values using a bayesian-allele-frequency approach. *J Genet Geneal*, 2:34–39, 2006.
- [3] K.P. Burnham and D.R. Anderson. *Model selection and multimodel inference: a practical information-theoretic approach*. Springer Verlag, 2002.
- [4] M.A. DePristo, E. Banks, R. Poplin, K.V. Garimella, J.R. Maguire, C. Hartl, A.A. Philippakis, G. Del Angel, M.A. Rivas, M. Hanna, et al. A framework for variation discovery and genotyping using next-generation dna sequencing data. *Nature genetics*, 43(5):491–498, 2011.
- [5] R.M. Durbin, D.L. Altshuler, G.R. Abecasis, D.R. Bentley, A. Chakravarti, A.G. Clark, F.S. Collins, F.M. De La Vega, P. Donnelly, M. Egholm, et al. A map of human genome variation from population-scale sequencing. *Nature*, 467(7319):1061–1073, 2010.

- [6] M.V. Han and C.M. Zmasek. phyloxml: Xml for evolutionary biology and comparative genomics. *BMC bioinformatics*, 10(1):356, 2009.
- [7] T.M. Karafet, F.L. Mendez, M.B. Meilerman, P.A. Underhill, S.L. Zegura, and M.F. Hammer. New binary polymorphisms reshape and increase resolution of the human y chromosomal haplogroup tree. *Genome Research*, 18(5):830, 2008.
- [8] H. Li, B. Handsaker, A. Wysoker, T. Fennell, J. Ruan, N. Homer, G. Marth, G. Abecasis, and R. Durbin. The sequence alignment/map format and samtools. *Bioinformatics*, 25(16):2078, 2009.
- [9] Y. Li, C. Sidore, H.M. Kang, M. Boehnke, and G.R. Abecasis. Low-coverage sequencing: Implications for design of complex trait association studies. *Genome Research*, 2011.
- [10] S. Purcell, B. Neale, K. Todd-Brown, L. Thomas, M. A. Ferreira, D. Bender, J. Maller, P. Sklar, P. I. de Bakker, M. J. Daly, and P. C. Sham. PLINK: a tool set for whole-genome association and population-based linkage analyses. *Am. J. Hum. Genet.*, 81:559–575, Sep 2007.
- [11] E. Sezgin, J.M. Lind, S. Shrestha, S. Hendrickson, J.J. Goedert, S. Donfield, G.D. Kirk, J.P. Phair, J.L. Troyer, S.J. OBrien, et al. Association of y chromosome haplogroup i with hiv progression, and haart outcome. *Human genetics*, 125(3):281–294, 2009.
- [12] P.A. Underhill, N.M. Myres, S. Rootsi, M. Metspalu, L.A. Zhivotovsky, R.J. King, A.A. Lin, C.E.T. Chow, O. Semino, V. Battaglia, et al. Separating the post-glacial coancestry of european and asian y chromosomes within haplogroup r1a. *European journal of human genetics*, 18(4):479–484, 2009.